Machine Learning with Memristors via Thermodynamic RAM

Timothy W. Molter
Knowm Inc., Santa Fe, NM, USA
Email: tim@knowm.org

M. Alexander Nugent
Knowm Inc., Santa Fe, NM, USA
Email: alex@knowm.org

arXiv:1608.04105v1 [cs.ET] 14 Aug 2016

***Abstract*—Thermodynamic RAM (kT-RAM) is a neuromemristive co-processor design based on the theory of AHaH Computing and implemented via CMOS and memristors. The co-processor is a 2-D array of differential memristor pairs (synapses) that can be selectively coupled together (neurons) via the digital bit addressing of the underlying CMOS RAM circuitry. The chip is designed to plug into existing digital computers and be interacted with via a simple instruction set. Anti-Hebbian and Hebbian (AHaH) computing forms the theoretical framework from which a nature-inspired type of computing architecture is built where, unlike von Neumann architectures, memory and processor are physically combined for synaptic operations. Through exploitation of AHaH attractor states, memristor-based circuits converge to attractor basins that represents machine learning solutions such as unsupervised feature learning, supervised classification and anomaly detection. Because kT-RAM eliminates the need to shuttle bits back and forth between memory and processor and can operate at very low voltage levels, it can significantly surpass CPU, GPU, and FPGA performance for synaptic integration and learning operations. Here, we present a memristor technology developed for use in kT-RAM, in particular bi-directional incremental adaptation of conductance via short low-voltage (<1.0 V, <1.0 μS) pulses.**

I. Introduction

Machine learning applications span a very diverse landscape. Some areas include motor control, combinatorial search and optimization, clustering, prediction, anomaly detection, classification, regression, natural language processing, planning and inference. A common thread is that a system learns the patterns and structure of the data in its environment by building a model, and then uses that model to make predictions of subsequent events and take actions. The models that emerge contain hundreds to trillions of continuously adaptive parameters. The human brain contain on the order of 10^{15} adaptive synapses. How the adaptive variables are exactly implemented in an algorithm varies, and established methods include support vector machines, decision trees, and artificial neural networks to name a few. Intuition tells us learning and modeling the environment is a valid approach in general, and the biological brain also appears to operate in this manner.

An unfortunate limitation when combining large-scale adaptive learning with computing is the memory-processing duality. In a traditional digital computer, calculations over adaptive variables must necessarily be performed in different physical locations than where the variables' memory resides, and this can often be a significant distance. The power required to adapt parameters grows impractically large as the number of parameters increases owing to the tremendous energy consumed shuttling digital bits back and forth between memory and processor. In a biological brain (and all of Nature), the processor and memory are the same physical substrate and computations and memory adaptations are performed in parallel. A truly low-power solution to machine learning occurs when the memory-processor distance goes to zero, and this can be achieved through intrinsically adaptive hardware such as memristors.

kT-RAM was designed with a few key constraints in mind, which are different than existing approaches: (1) must result in a hardware solution where memory and computation are combined, (2) must not only solve a single problem but provide a generic solution to most machine learning applications and (3) must be simple to fabricate with existing or emerging manufacturing technology. This initial motivation led us to create a theoretical framework for a neuromorphic processor satisfying the above constraints, which we call AHaH computing [1]. Independent works such as the Local Activity Principle [2] and Thermodynamics of Economics Modeling [3], while more general theoretical frameworks compared to AHaH Computing, are based on roughly the same ideas of homogeneous matter spontaneously forming complex structure in order to maximize energy dissipation (4th Law of Thermodynamics).

Whereas previous publications present both AHaH Computing and Thermodynamic RAM in detail [1], [4], [5], here we will focus on the memristors we have designed and plan to use in the first kT-RAM prototype chips.

II. Results and Discussion

We have recently developed and made commercially-available memristors with bi-directional incremental learning capability. This advancement opens the gateway to extremely efficient and powerful machine learning and artificial intelligence applications. The memristor devices come in three variants: W, Sn and Cr, which refers to the metal introduced in the active layer during fabrication resulting in different electrical characteristics. The memristor devices operate primarily through the mechanism of electric field induced generation and movement of metal ions through a multilayer chalcogenide material stack. The memristor devices are fabricated with a layer of metal that is easily oxidizable, located near one electrode. When a voltage is applied across the device with the

more positive potential on the electrode near this metal layer, the metal is oxidized to form ions. Once formed, the ions move through the device towards the lower potential electrode. The ions move through a layer of amorphous chalcogenide material (the active layer) to reach the lower potential electrode where they are reduced to their metallic form and eventually form a conductive pathway between both electrodes that spans the active material layer, lowering the device resistance. Reversing the direction of the applied potential causes the conductive channel to dissolve and the device resistance to increase. The devices are bipolar, cycling between high and low resistance values by switching the polarity of the applied potential. The resistance is related at any time to the amount of metal located within the active layer. A secondary phase change mechanism can be used beforehand to initialize the memristor devices into the desired high and low conductance range.

For bi-directional pulsed incremental learning applications, such as in kT-RAM, the primary mode of operation is the fine-grained metal ion transport within the device (Figure 1). Given one or more low-voltage pulses ($<$1.0 V, $<$1.0 μS), the memristor's conductance can be "nudged" into either a higher or lower value, representing a synaptic weight for example. In kT-RAM synapses are realized with two serially-connected memristors and used as a voltage divider where the output voltage is the synapse's weight. A read pulse across the memristor pair provides a "weight" and drives the synapse towards a forgetful state, while the subsequent write pulse across one of the memristors provides a learning signal to the synapse.

Using a memristor model called the "Metastable Switch Model" [1] we have developed a simulation platform to run machine learning applications and benchmarks on top of emulated kT-RAM circuits. To date we have demonstrated state-of-the-art benchmark performance on a handful of standard machine learning classification benchmarks including the popular MNIST hand-written digits dataset (1.0% error rate). Additionally, we have demonstrated real-time clustering, time-series signal prediction, combinatorial optimization, anomaly detection, unsupervised feature learning and reinforcement learning. Demonstrations have been previously published in "AHaH Computing – From Metastable Switches to Attractors to Machine Learning" [1] including companion open source code (https://github.com/timmolter/ahah), a web article series titled Machine Learning Capabilities with Thermodynamic RAM and the Knowm API (http://knowm.org/machine-learning-capabilities-with-thermodynamic-ram-and-the-knowm-api/), and a play list on YouTube called Emulated Thermodynamic RAM Demonstrations (https://www.youtube.com/playlist?list=PLlDmU5euIpk0qiErPtBhtrtL3_53jtnL9). Recent advancements in deep learning from various groups have further improved the scores on ML benchmarks, and we are currently developing our own deep learning capabilities on emulated kT-RAM and tackling more challenging image-recognition benchmarks. To date we have successfully produced the memristors we need that exhibit bi-directional incremental learning as presented here. The next steps are to build the first AHaH Node prototype circuits, followed by fully-functional Thermodynamic RAM.

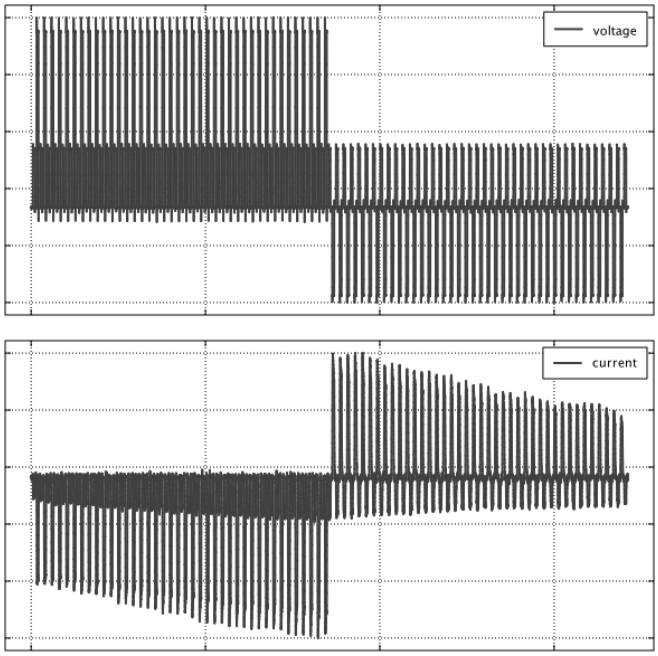

Fig. 1. A Tungsten Ag-chalcogenide memristor device exhibits bi-directional incremental learning behavior in response to 50 ns voltage pulses; useful for analog-based learning architectures such as Thermodynamic RAM. The voltage plot shows asymmetric forward and reverse activation pulses as well as lower-magnitude read pulses. The current plot shows the corresponding incremental change in conductance of the device occurring in both increasing and decreasing directions.

ACKNOWLEDGMENT

The authors would like to thank the Air Force Research Labs in Rome, NY for their support under the SBIR/STTR programs AF10-BT31, AF121-049. The authors would like to thank Kristy A. Campbell from Boise State University for graciously providing us with memristor device data.

REFERENCES

[1] M. A. Nugent and M. T. W, "Ahah computing-from metastable switches to attractors to machine learning," *PLoS ONE*, vol. 9, p. e85175, 02 2014.
[2] K. Mainzer and L. Chua, "The cause of complexity and symmetry breaking," 2013.
[3] S. E. Jørgensen, *Evolutionary Essays:: A Thermodynamic Interpretation of the Evolution*. Elsevier, 2011.
[4] M. A. Nugent and M. T. W, "Thermodynamic-ram technology stack," *arXiv preprint arXiv:1406.5633*, 2014.
[5] ——, "Cortical processing with thermodynamic-ram," *arXiv preprint arXiv:1408.3215*, 2014.